<u>First Hit</u>   <u>Fwd Refs</u>
End of Result Set

☐ | Generate Collection | | Print |
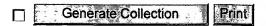
L25: Entry 1 of 1   File: USPT   Feb 6, 2001

DOCUMENT-IDENTIFIER: US 6185598 B1
TITLE: Optimized network resource location

<u>Brief Summary Text</u> (4):
The advent of global computer networks, such as the Internet, have led to entirely
new and different ways to obtain information. A <u>user</u> of the Internet can now access
information from anywhere in the world, with no regard for the actual location of
either the <u>user</u> or the information. A <u>user</u> can obtain information simply by knowing
a network address for the information and providing that address to an appropriate
application program such as a network browser.

<u>Brief Summary Text</u> (13):
Assuming that client protocols cannot easily change, there are two major problems
in the deployment of replicated resources. The first is how to select which copy of
the resource to use. That is, when a request for a resource is made to a single
server, how should the choice of a replica of the server (or of that data) be made.
We <u>call</u> this problem the "rendezvous problem". There are a number of ways to get
clients to rendezvous at distant mirror servers. These technologies, like load
balancers, must route a request to an appropriate server, but unlike load
balancers, they take network performance and topology into account in making the
determination.

<u>Brief Summary Text</u> (14):
A number of companies offer products which improve network performance by
prioritizing and filtering network traffic. Proxy caches provide a way for client
aggregators to reduce network resource consumption by storing copies of popular
resources close to the end <u>users</u>. A client aggregator is an Internet service
provider or other organization that brings a large number of clients operating
browsers to the Internet. Client aggregators may use proxy caches to reduce the
bandwidth required to serve web content to these browsers. However, traditional
proxy caches are operated on behalf of Web clients rather than Web publishers.

<u>Detailed Description Text</u> (7):
Client 106 is a processor requesting resources from origin server 102 on behalf of
an end <u>user</u>. The client 106 is typically a <u>user</u> agent (e.g., a Web browser such as
Netscape Communications Corporation's Navigator.TM.) or a proxy for a <u>user</u> agent.
Components other than the reflector 108 and the repeaters 104a, 104b, etc., may be
implemented using commonly available software programs. In particular, this
invention works with any HTTP client (e.g., a Web browser), proxy cache, and Web
server. In addition, the reflector 108 might be fully integrated into the data
server 112 (for instance, in a Web Server). These components might be loosely
integrated based on the use of extension mechanisms (such as so-called add-in
modules) or tightly integrated by modifying the service component specifically to
support the repeaters.

<u>Detailed Description Text</u> (17):
A hypertext document may contain any number of links to other documents, and each
of those other documents may be on a different server in a different part of the

h   e b   b g e e e f   c   e   b   e   ge

world. For example, a document may contain links to documents in Russia, Africa, China and Australia. A user viewing that document at a particular client can follow any of the links transparently (i.e., without knowing where the document being linked to actually resides). Accordingly, the cost (in terms of time or money or resource allocation) of following one link versus another may be quite significant.

Detailed Description Text (21):
In order to access the file "F" (the resource) specified by the above URL, a program (e.g., a browser) running on a user's computer (i.e., a client computer) would have to first locate the computer (i.e., a server computer) specified by the host name. I.e., the program would have to locate the server "www.uspto.gov". To do this, it would access a Domain Name Server (DNS), providing the DNS with the host name ("www.uspto.gov"). The DNS acts as a kind of centralized directory for resolving addresses from names. If the DNS determines that there is a (remote server) computer corresponding to the name "www.upsto.gov", it will provide the program with an actual computer network address for that server computer. On the Internet this is called an Internet Protocol (or IP) address and it has the form "123.345.456.678". The program on the user's (client) computer would then use the actual address to access the remote (server) computer.

Detailed Description Text (24):
A1. A browser (e.g., Netscape's Navigator) at the client receives a resource identifier (i.e., a URL) from a user.

Detailed Description Text (32):
There are many variations of this basic model. For example, in one variation, instead of providing the client with the resource, the origin server can tell the client to re-request the resource by another name. To do so, in A7 the server 102 sends back to the client 106 a reply called a "REDIRECT" which contains a new URL indicating the other name. The client 106 then repeats the entire sequence, normally without any user intervention, this time requesting the resource identified by the new URL.

Detailed Description Text (286):
There are several cases where a resource may be cached so long as the origin server is consulted each time it is served. In one case, the request for the resource is attached to a so-called "cookie". The origin server must be presented with the cookie to record the request and determine whether the cached resource may be served or not. In another case, the request for the resource is attached to an authentication header (which identifies the requester with a user id and password). Each new request for the resource must be tested at the origin server to assure that the requester is authorized to access the resource.

h   e b   b g ee e f   c   e   b   e ge